ATOMIC HINS: ENTITY-ATTRIBUTE DUALITY FOR HETEROGENEOUS GRAPH MODELING

Anonymous authors

000

002 003 004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

032

034

035

037

040

041

042

043

044

046

047

048

050

051

052

Paper under double-blind review

ABSTRACT

Heterogeneous Information Networks (HINs) provide a powerful framework for modeling multi-typed entities and relations, typically defined under a fixed schema. Yet, most research assumes this structure is given, overlooking the fact that alternative designs can emphasize different aspects of the data and substantially influence downstream performance. As a theoretical foundation for such designs, we introduce the principle of entity-attribute duality: attributes can be atomized as entities with their associated relations, while entities can, in turn, serve as attributes of others. This principle motivates atomic HIN, a canonical representation that makes all modeling choices explicit and achieves maximal expressiveness. Building on this foundation, we propose a systematic framework for task-specific schema refinement. Within this framework, we demonstrate that widely used benchmarks correspond to heuristic refinements of the atomic HINoften far from optimal. Across eight datasets, refinement alone enables a simplified Relational GCN (sRGCN) to reach state-of-the-art performance on node- and link-level tasks, with further gains from advanced HGNNs. These results highlight schema design as a key dimension in heterogeneous graph modeling. By releasing the atomic HINs, searched schemas, and refinement framework, we enable principled benchmarking and open the way for future work on schema-aware learning, automated structure discovery, and next-generation HGNNs.

1 Introduction

Heterogeneous Information Networks (HINs) provide a powerful abstraction for modeling systems with multiple types of entities and relations. Such graphs naturally arise in bibliometrics, e-commerce, knowledge graphs, and social networks, where diverse node and edge types yield rich semantics. To leverage these structures, Heterogeneous Graph Neural Networks (HGNNs) extend Graph Neural Networks (GNNs) with type-aware message passing across heterogeneous schemas.

Despite this progress, most research relies on a few benchmark HINs with manually specified schemas, often chosen heuristically. In practice, however, multiple valid schemas can be derived from the same data. For instance, the IMDB benchmark is constructed from a single movie table¹: columns such as *actor* and *director* are processed to be entities, while others (e.g., keyword, language, country) remain attributes (Figure 1b). In some variants, *keywords* are treated as entities rather than mere attributes (Figure 1c), while additional unexplored variants are equally possible (Figure 1d).

The ambiguity in schema design has received limited attention. Some recent efforts (Fey et al., 2024) explore schema construction from relational databases, but they still rely on database-specific design choices, each of which corresponds to different schema variants. Consequently, the broader problem of designing HIN schemas remains open.

We introduce *atomic HIN*, grounded in the principle of *entity-attribute duality*: attributes can be atomized as entities, and entities can in turn serve as attributes of others. This duality maximizes expressiveness by making all schema choices explicit, but also increases modeling complexity.

https://www.kaggle.com/datasets/karrrimba/movie-metadatacsv

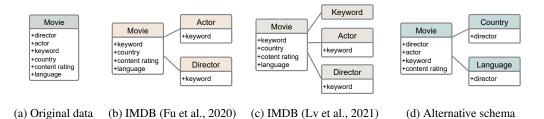


Figure 1: Illustration of multiple heterogeneous schemas derived from the same IMDB source table.

To address this, we propose *task-specific schema refinement*, which selects or discards node types (entities) and edge types (relations) to tailor the schema. Node-type selection determines which entities are assigned unique identities, while edge-type selection removes relations of limited utility, simplifying both the schema and the model. From this perspective, widely used benchmark schemas are merely heuristic node- and edge-type selections, as essentially all schema variants can be viewed as refinements of the atomic form.

Empirically, pairing the atomic form with sRGCN—a simplified RGCN with stronger parameter sharing—achieves state-of-the-art results on eight datasets spanning node classification and link prediction. On node classification, Macro-F1 improves by up to 6.2% over recent advanced HGNNs, and on link prediction, ROC-AUC increases by an average of 4.9%. Our systematic search further shows that benchmark schemas often diverge from optimal ones, highlighting the effectiveness of *atomic HINs* with schema refinement. Moreover, schemas refined with sRGCN generalize well to advanced HGNNs, with subsequent refinement under stronger models yielding additional gains.

We summarize our contributions as follows:

- We introduce the principle of *entity-attribute duality*, which motivates the *atomic HIN*, a canonical representation that makes all schema choices explicit and achieves maximal expressiveness.
- To manage the resulting complexity, we propose task-specific schema refinement, formulated as
 node- and edge-type selection. Existing benchmark schemas emerge as particular refinements
 of the atomic form, but our systematic search shows that such heuristic choices often diverge
 from optimal ones, underscoring the central role of schema design in HGNN development and
 evaluation.
- Extensive experiments on eight benchmarks demonstrate that *atomic HINs* with refined schemas consistently improve HGNN performance. Even a simplified RGCN (sRGCN) achieves state-of-the-art results, while advanced HGNNs obtain further improvements. Analysis of the search results further confirms the importance of *atomic HINs* and schema refinement.

2 Preliminaries

Attributed Heterogeneous Information Networks

An undirected HIN is a graph $\mathcal{G}=(\mathcal{V},\mathcal{E})$ consisting of a node set \mathcal{V} and an edge set \mathcal{E} . Each node $v\in\mathcal{V}$ and edge $e\in\mathcal{E}$ is assigned a type through mappings $\phi:\mathcal{V}\to\mathcal{T}$ and $\psi:\mathcal{E}\to\mathcal{R}$, where $\mathcal{T}=\left\{\tau\right\}_{\tau=1}^{|\mathcal{T}|}$ and $\mathcal{R}=\left\{r\right\}_{r=1}^{|\mathcal{R}|}$ denote the sets of node types (entities) and edge types (relations). Edges are represented by adjacency matrices $\left\{A_r\in\mathbb{R}^{|\mathcal{V}|\times|\mathcal{V}|}\right\}_{r=1}^{|\mathcal{R}|}$, where A_r corresponds to edges with relation $r\left\{e\in\mathcal{E}\mid\psi(e)=r\right\}$. When $|\mathcal{T}|=|\mathcal{R}|=1$, the network \mathcal{G} reduces to a homogeneous graph.

An attributed HIN is further associated with an attribute set \mathcal{F} , where each $f \in \mathcal{F}$ is assigned to an owner node type via $\zeta: \mathcal{F} \to \mathcal{T}$. Attributes are then represented as feature matrices $\{ \boldsymbol{X}_f \in \mathbb{R}^{|\mathcal{V}_{\zeta(f)}| \times d_f} \}_f$, where $\mathcal{V}_{\zeta(f)} = \{ v \in \mathcal{V} \mid \phi(v) = \zeta(f) \}$ and d_f is the dimension of f.

Spectral heterogeneous graph convolution (SHGC). Defferrard et al. (2016) showed that spectral filters on homogeneous graphs can be parameterized as L-order polynomials of the Laplacian eigenvalues. Butler et al. (2023) extended the formulation to heterogeneous graphs, introducing spectral filters based on non-commutative polynomials over relation-specific shift operators. For an

input signal $x \in \mathbb{R}^{|\mathcal{V}|}$, the filter is defined as

$$\boldsymbol{H}(\boldsymbol{I}, \boldsymbol{S}_1, \dots, \boldsymbol{S}_{|\mathcal{R}|}; \boldsymbol{\Theta}) \boldsymbol{x} = \theta_0 \boldsymbol{I} \boldsymbol{x} + \sum_{\ell=1}^{L} \sum_{r_1, \dots, r_\ell} \theta_{r_1, \dots, r_\ell} (\boldsymbol{S}_{r_1} \cdots \boldsymbol{S}_{r_\ell}) \boldsymbol{x}, \tag{1}$$

where S_r is the shift operator for relation r. A common choice is the row-normalized adjacency $\tilde{A}_r = D_r^{-1} A_r$, where $D_r[i,i] = \sum_{j=1}^{|\mathcal{V}|} |A_r[i,j]|$ is the degree matrix of A_r . The learnable parameters are collected in $\Theta = \{\theta_0\} \cup \{\theta_{r_1,\dots,r_\ell}\}_{r_1,\dots,r_\ell}$, where each coefficient $\theta_0, \theta_{r_1,\dots,r_\ell} \in \mathbb{R}$. In this paper, we adopt SHGC as a general formulation of HGNNs.

3 METHODOLOGY

3.1 From Attribute Atomization to Atomic HINs

Constructing graph structure from attributes is a long-standing but often implicit technique in HIN schema design. In practice, this is typically applied during preprocessing to attributes that can be represented by one-hot or multi-hot encodings, thereby introducing additional node and edge types. A widely used example is the IMDB dataset (Figure 1). Formally, the process can be applied to each attribute independently. For a given attribute, the procedure is defined as follows:

Definition 3.1 (Attribute Atomization). Given an attribute $f \in \mathcal{F}$ of a HIN \mathcal{G} with feature matrix X_f , atomizing f produces an augmented HIN $\mathcal{G}' = (\mathcal{V} \cup \mathcal{U}, \mathcal{E} \cup \mathcal{E}_f)$, where $\mathcal{U} = \{u_1, \dots, u_{d_f}\}$ is the set of new nodes with cardinality corresponding to attribute dimensions and

$$\mathcal{E}_f = \{ (v_i, u_j) \mid X_f[i, j] \neq 0, \ 1 \leq i \leq |\mathcal{V}_\tau|, \ 1 \leq j \leq d_f \}.$$

Here v_i denotes the i-th node of $\mathcal{V}_{\zeta(f)}$ and u_j is the j-th node of the induced set \mathcal{U} . Each edge (v_i,u_j) is weighted by $X_f[i,j]$. This introduces a new node type τ' and a new edge type r', yielding $\mathcal{T}' = \mathcal{T} \cup \{\tau'\}$ with $\phi'(v) = \tau'$ if $v \in \mathcal{U}$ and $\phi'(v) = \phi(v)$ otherwise; and $\mathcal{R}' = \mathcal{R} \cup \{r'\}$ with $\phi'(e) = r'$ if $e \in \mathcal{E}_f$ and $\phi'(e) = \phi(e)$ otherwise.

In this way, atomization replaces nonzero entries of X_f with explicit edges to induced nodes, thereby converting attributes into structure and relation. This enables HGNNs to exploit structural dependencies and relational patterns (e.g., metapaths (Sun et al., 2011)). The common practice of constructing edges by ID matching across relational tables (Fey et al., 2024) is a special case where X_f is sparse and binary.

Applying atomization to all attributes yields the *atomic HIN*, a canonical representation in which all information is expressed structurally, achieving maximal expressiveness and making explicit all modeling choices. This representation serves as the foundation for our next step: task-specific schema refinement.

3.2 SCHEMA REFINEMENT VIA NODE- AND EDGE-TYPE SELECTION

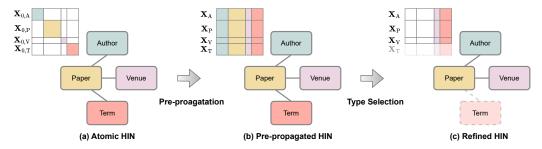


Figure 2: Toy example of schema refinement on *atomic HINs*: pre-propagation ensures independence between node-type and edge-type selection.

We then propose the schema refinement problem, a binary selection problem over node types (entities) and edge types (relations), aiming to harness the expressiveness of *atomic HINs* while controlling complexity. To this end, we first define two basic operations:

- Node-type selection. For each node type $\tau \in \{1, \dots, |\mathcal{T}|\}$, a binary parameter $\beta_{\tau} \in \{0, 1\}$ determines whether nodes of this type are assigned a unique identity and thus contribute as attributes to the HIN.
- Edge-type selection. For each relation $r \in \{1, ..., |\mathcal{R}|\}$, a binary parameter $\alpha_r \in \{0, 1\}$ specifies whether the relation is retained and modeled in the schema.

Under the SHGC framework (Equation 1), which we use as a general form of HGNNs, refinement on an *atomic HIN* is expressed as:

$$\bar{Z} = H(I, \alpha_1 S_1, \dots, \alpha_{|\mathcal{R}|} S_{|\mathcal{R}|}; \Theta) X(\beta_1, \dots, \beta_{|\mathcal{T}|}),$$
(2)

where $H(\cdot)$ denotes the SHGC filter, and $X(\beta_1, \dots, \beta_{|\mathcal{T}|}) \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the feature matrix after node-type selection.

The edge-type selectors α_r directly control which relations and edges participate in message passing. Setting $\alpha_r=0$ removes relation r and its edges entirely, equivalent in practice to dropping that relation from the constructed HIN. This makes refinement compatible with any HGNN in a plugand-play manner, without requiring architectural modifications.

Node-type selection specifies which node types are assigned unique identities, i.e., learnable embeddings. A straightforward construction is to assign the identity matrix as node features for the selected node types:

$$\boldsymbol{X}_0(\beta_1, \dots, \beta_{|\mathcal{T}|}) = \beta_1 \hat{\boldsymbol{I}}_1 + \beta_2 \hat{\boldsymbol{I}}_2 + \dots + \beta_{|\mathcal{T}|} \hat{\boldsymbol{I}}_{|\mathcal{T}|}, \tag{3}$$

where $\hat{I}_{\tau} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the type-specific identity matrix, with $\hat{I}_{\tau}[i,i] = 1$ if the i-th node $v_i \in \mathcal{V}$ has type $\phi(v_i) = \tau$, and 0 otherwise. When all types are selected, $X_0(1,\ldots,1) = I_{|\mathcal{V}|}$. The motivation is to learn embeddings only for informative node types, thereby reducing parameters and mitigating overfitting.

However, this naive construction can induce dependencies.

Definition 3.2 (Dependency relative to a node type). For a node type $\tau_i \in \{1, ..., |\mathcal{T}|\}$, we say that τ_i has a dependency on τ_j if $\hat{I}_{\tau_i}\bar{Z}_0\hat{I}_{\tau_j} = 0$ for some τ_j with $\beta_{\tau_j} = 1$, where $\bar{Z}_0 = H(\cdot)X_0(\cdot)$ is the output of SHGC on the naive feature initialization.

Dependencies arise because pruning certain edge types can disconnect selected nodes from the rest of the graph. For example, in Figure 2, if all edges incident to *term* nodes are removed, their embeddings become isolated and cannot contribute to downstream predictions. Thus, although edge-type pruning improves efficiency and compatibility, it may inadvertently invalidate node-type selections by cutting off identity propagation.

To mitigate this issue, we introduce a pre-propagation feature initialization, which ensures that embeddings of selected node types remain accessible regardless of subsequent edge-type choices.

3.3 Pre-propagation Feature Initialization

To address the dependency issue in Definition 3.2, we introduce a pre-propagation feature initialization. The idea is straightforward: prior to refinement, each selected node type propagates its identity once to other node types, ensuring that its signal remains accessible even if incident edge types are subsequently removed. Formally, we define pre-propagated features as

$$\boldsymbol{X}(\beta_1, \dots, \beta_{|\mathcal{T}|}) = \left(\boldsymbol{I} + \sum_{\tau_i \neq \tau_j} \tilde{\boldsymbol{A}}_{\langle \tau_i, *, \tau_j \rangle} \right) \boldsymbol{X}_0(\beta_1, \dots, \beta_{|\mathcal{T}|}), \tag{4}$$

where $\tilde{A}_{\langle \tau_i, *, \tau_j \rangle}$ denotes the adjacency of the shortest path from type τ_j to τ_i , defined as the shortest product of adjacency matrices satisfying $\hat{I}_{\tau_i} \tilde{A}_{\langle \tau_i, *, \tau_j \rangle} \hat{I}_{\tau_j} \neq 0$. We provide details for the implementation of this pre-propagation step in Appendix B.5.

This initialization effectively distributes each identity embedding beyond its original source node, ensuring its availability across node types and thereby eliminating hidden dependencies.

Lemma 3.1 (Independence of Selections). With pre-propagation, for any node type τ_i and any selected type τ_j with $\beta_{\tau_j} = 1$, $\hat{I}_{\tau_i} \bar{Z} \hat{I}_{\tau_j} \neq 0$. Hence, node-type selection is independent of edge-type selection: even if all edges incident to τ_j are removed, no dependency arises.

Lemma 3.2 (Neutrality of Pre-Propagation). Consider an SHGC with row-normalized adjacencies $S_r = \tilde{A}_r$ as shift operators and convolution order L sufficiently large. If raw identity features $X_0(1,\ldots,1) = I$ are replaced with pre-propagated features $X(1,\ldots,1)$, the effect is merely a reparameterization of the filter coefficients $\{\theta_{r_1,\ldots,r_\ell}\}_{r_1,\ldots,r_\ell}$. Hence, pre-propagation does not alter the expressive power of the model.

Proofs are deferred to Appendix B.

In summary, pre-propagation resolves node dependencies (Lemma 3.1) without reducing expressive power (Lemma 3.2), ensuring that edge types can be pruned without invalidating node type selections. We illustrate this with row-normalized adjacency, while in practice we adopt a generalized union via element-wise maximum (Lv et al., 2021) (see Appendix B.4). Pre-propagation thus provides a principled basis for schema refinement without manual adjustments for each choice.

3.4 Systematic Search for Schema Refinement

Schema refinement reduces to a binary selection over node and edge types, but the search space of $2^{|\mathcal{R}|+|\mathcal{T}|}$ candidates is prohibitively large. Moreover, the space is highly skewed: retaining more edge types generally increases expressiveness (Lemma 3.3), while sparse or high-cardinality node types often introduce excessive parameters and risk overfitting. Hence, naive grid or random search is ineffective.

To address this, we formulate schema refinement as a hyperparameter optimization problem and adopt a genetic algorithm (GA)—based search strategy. While GAs have been widely studied for binary optimization, to our knowledge, they have not been applied to HIN schema refinement. Our formulation naturally enables this approach, providing a practical and effective means to explore the skewed binary space and achieve near-optimal solutions within reasonable budgets (Deb et al., 2002; Katoch et al., 2021). We initialize the population with the vanilla schema and jointly optimize schema parameters with model depth L, consistent with Lemma 3.2. Full algorithmic details are provided in Appendix A.4.

3.5 HGNNs for Atomic HINs

Atomizing attributes enhances expressiveness by establishing additional structure and relations, enabling HGNNs to capture richer semantics. With all information represented structurally, assigning unique identities to entities allows them to act as attribute signals for others. In particular, the new relations introduced through attribute atomization enlarge the filter space, as formalized below:

Lemma 3.3 (Attribute Atomization Enlarges Filter Space). Let \mathcal{R} be the original set of edge types, and let $\mathcal{R}' \supset \mathcal{R}$ denote the set obtained after attribute atomization. Under the SHGC formulation, the space of heterogeneous filters spanned by

$$\{I\} \cup \{S_{r_1} \cdots S_{r_\ell} \mid r_1, \dots, r_\ell \in \mathcal{R}', 1 \le \ell \le L\}$$

is strictly larger than the corresponding space defined with \mathcal{R} . Hence, attribute atomization strictly enlarges the filter space by converting attributes into entities with associated relations.

To connect with existing HGNNs and illustrate the broad utility of *atomic HINs*, we observe that many architectures can be viewed as instances or approximations of SHGC.

Proposition 3.1 (RGCN (Schlichtkrull et al., 2018) as a First-Order Approximation of SHGC). RGCN can be expressed as a first-order approximation of SHGC under row-normalized adjacencies $S_r = \tilde{A}_r$, with filter coefficient matrices $\{W_0\} \cup \{W_r^{(\ell)}\}_{r,\ell}$.

Proposition 3.2 (GTN (Yun et al., 2019) as a First-Order Approximation of SHGC). GTN can be expressed as a first-order approximation of SHGC under row-normalized adjacencies, with scalar filter coefficients $\{\theta_0\} \cup \{\theta_r^{(\ell)}\}_{r,\ell}$.

These cases show that widely used HGNNs, despite differing in interpretation, are first-order SHGC with distinct parameter-sharing schemes (e.g., GTN enforces $\boldsymbol{W}_r^{(\ell)} = \theta_r^{(\ell)} \boldsymbol{I}$). Detailed derivations, along with extensions to higher-order variants such as SeHGNN (Yang et al., 2023) and PSHGCN (He et al., 2024), are provided in Appendix C.

On *atomic HINs*, SHGC particularly favors architectures with stronger parameter sharing (e.g., GTN-style), since all inputs reduce to unique identity embeddings. In this setting, heavy feature transformations or MLP layers become largely redundant—consistent with the empirical findings of He et al. (2020) that shallow parameterization can outperform deeper transformations when natural features are absent.

Motivated by this perspective, we introduce a simplified variant of RGCN, denoted *sRGCN*. It preserves the RGCN structure but replaces feature transformation matrices with relation-specific scalars,

$$\boldsymbol{W}_{r}^{(\ell)} = \theta_{r}^{(\ell)} \boldsymbol{I}.$$

This design yields a minimal yet effective baseline, naturally aligned with the requirements of *atomic HINs*. Full update rules and implementation details are provided in Appendix D.

4 EXPERIMENTS

In this section, we conduct experiments to address the following research questions:

- **RQ1:** Do *atomic HINs* improve performance over benchmark schemas and advanced HGNNs?
- **RQ2:** How does the *entity-attribute duality* in the atomic view benefit HGNNs?
- RO3: Is schema refinement truly necessary? How does it affect node-level and link-level tasks?
- **RQ4:** How large is the performance gain across schema variants when using the same HGNN? Do schemas refined on sRGCN generalize across HGNNs, and can they be further improved through subsequent refinement?
- **RQ5:** How efficient is schema refinement when optimized through a genetic algorithm?

4.1 EXPERIMENT SETUP

Datasets. We evaluate on eight heterogeneous benchmarks drawn from diverse domains, including bibliometrics, e-commerce, knowledge graphs, social networks, and biomedicine. Dataset statistics and corresponding vanilla schemas are listed in Table 4 and Table 2, respectively. Full dataset descriptions are provided in Appendix A.1. For all datasets, we perform attribute atomization on every available attribute. When datasets include initial embeddings (e.g., pretrained language-model (PLM) embeddings), we treat them as numerical attributes and likewise atomize them into feature nodes. For OGBN-MAG, following Yang et al. (2023), we initialize large-scale node types with 256-dimensional random embeddings to approximate learnable ID embeddings.

Evaluation Setting. We follow the experimental protocols of each dataset, as specified by its benchmark or commonly adopted in the literature. For schema refinement, we employ a GA algorithm with 1024 candidates, and subsequently fine-tune HGNN hyperparameters on the derived optimal schema with 256 trials. Full details of experimental setup, baseline implementations, and hyperparameter space are provided in Appendix A.

4.2 Performance on Atomic HINs (RQ1)

We first evaluate the effectiveness of *atomic HINs* using sRGCN across eight benchmarks for node classification and link prediction. We compare against recent state-of-the-art HGNNs under their vanilla schemas. As shown in Table 1, sRGCN on refined atomic schemas consistently outperforms advanced HGNNs. Gains are more significant on datasets with rich attributes or complex schemas (e.g., IMDB, Amazon, Freebase), where attribute atomization increases expressiveness while refinement tailors complexity. ACM is already in atomic form with a simple schema, leaving little room for improvement. Yet, LastFM benefits substantially from refinement despite also being in atomic form with an even simpler schema. On the large-scale OGBN-MAG, improvements are mainly driven by relations induced from PLM embeddings—atomized as relation nodes—together with schema refinement, showing that even numerical attributes can yield useful relational patterns.

Overall, Macro-F1 improves by up to 6.2% for node classification and ROC-AUC by an average of 4.9% for link prediction over the strongest baselines.

Table 1: Performance comparison on different datasets. Top panel reports **node classification** results; bottom panel reports **link prediction** results. Best results are in **bold**, second-best are <u>underlined</u>. Statistical significance is marked by † (p < 0.001) and ‡ (p < 0.01). OOM indicates out-of-memory.

	IM	DB	Free	base	DB	SLP	A(CM	OGBN-MAG		
HGNN	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Acc. (Val)	Acc. (Test)	
RGCN	58.85±0.26	62.05±0.15	46.78±0.77	58.33±1.57	91.52±0.50	92.07±0.50	91.55±0.74	91.41±0.75	48.35±0.36	47.37±0.48	
HGT	63.00 ± 1.19	67.20 ± 0.57	29.28 ± 2.52	60.51 ± 1.16	93.01 ± 0.23	93.49 ± 0.25	91.12 ± 0.76	91.00 ± 0.76	49.89 ± 0.47	49.27 ± 0.61	
SimpleHGN	63.53 ± 1.36	67.36 ± 0.57	47.72 ± 1.48	66.29 ± 0.45	94.01 ± 0.24	94.46 ± 0.22	93.42 ± 0.44	93.35 ± 0.45	OOM	OOM	
HINormer	64.65 ± 0.53	67.83 ± 0.34	52.18 ± 0.39	64.92 ± 0.43	94.57 ± 0.23	94.94 ± 0.21	93.91 ± 0.42	93.83 ± 0.45	OOM	OOM	
SeHGNN	66.63 ± 0.34	68.21 ± 0.32	50.71 ± 0.44	63.41 ± 0.47	94.86 ± 0.14	95.24 ± 0.13	93.95 ± 0.48	93.87 ± 0.50	55.95 ± 0.11	53.99 ± 0.18	
SlotGAT	64.05 ± 0.60	68.64 ± 0.33	49.68 ± 1.97	66.83 ± 0.30	94.95 ± 0.20	95.31 ± 0.19	93.99 ± 0.23	94.06 ± 0.22	OOM	OOM	
PSHGCN	67.10 ± 0.60	69.79 ± 0.52	40.01 ± 8.26	$\overline{62.70\pm0.77}$	95.27 ± 0.13	95.61 ± 0.12	94.35 ± 0.23	94.27 ± 0.23	56.16 ± 0.21	54.57 ± 0.16	
sRGCN _{Atomic}	$68.97 \pm 0.09^{\ddagger}$	$71.20\pm0.17^{\dagger}$	$55.40 \pm 1.25^{\dagger}$	67.32 ± 0.66	$95.55 \pm 0.13^{\ddagger}$	$95.85 \pm 0.12^{\ddagger}$	94.36 ± 0.22	94.29 ± 0.22	$57.35 \pm 0.12^{\dagger}$	$55.21 \pm 0.23^{\dagger}$	

	Am	azon	Las	tFM	PubMed			
HGNN	ROC-AUC	MRR	ROC-AUC	MRR	ROC-AUC	MRR		
RGCN	86.34±0.24	93.92±0.16	57.21±0.09	77.68±0.17	84.62±0.33	94.27±0.51		
HGT	88.26 ± 2.06	93.87 ± 0.65	54.99 ± 0.28	74.96 ± 1.46	85.38 ± 1.20	94.98 ± 0.69		
SimpleHGN	93.40 ± 0.62	96.94 ± 0.29	67.59 ± 0.23	90.81 ± 0.32	85.48 ± 1.08	93.67 ± 1.06		
SeHGNN	91.67 ± 0.94	95.83 ± 0.58	66.59 ± 0.62	88.61 ± 1.25	85.86 ± 1.11	95.09 ± 0.74		
SlotGAT	95.17 ± 0.11	98.00 ± 0.09	70.33 ± 0.13	91.72 ± 0.50	88.07 ± 0.20	94.71 ± 0.33		
PSHGCN	94.12±0.58	$\overline{97.93\pm0.46}$	$\overline{69.25\pm0.63}$	91.19 ± 0.51	87.16 ± 1.89	95.01 ± 1.26		
sRGCN	97.85±0.07 [†]	99.26+0.05	$77.10\pm0.17^{\dagger}$	93.70±0 16 [†]	90.11+0.19	96.14+0.04		

Table 2: Schema refinement results on sRGCN. Edge types with <u>underline</u> are weighted, and with <u>double</u> <u>underline</u> are weighted and dense. Node type aliases are listed in Appendix A.1.1.

Dataset	Variant		Node types						Edge types											
		N	R	D	W	C	L	A	О	K	M	M-D	M-K	M-W	M-A	M-R	M-N	M-L	М-О	М-С
IMDB	Vanilla	✓	✓	-	✓	✓	✓	-	✓	√	-	✓	✓	-	✓	-	-	-	-	-
	Refined	✓	✓	✓	✓	✓	✓	-	-	-	-	✓	✓	✓	✓	✓	✓	-	-	-
		Bk	\mathbf{Or}	$\mathbf{B}\mathbf{s}$	Mu	Sp	Lo	Fi	Pe			Bk-Bk	Bk-Pe	Bk-Or	Bk-Bs		Pe-Sp	Bk-Fi	Mu-Bs	Pe-Or
Freebase	Vanilla	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓	✓	✓		✓	✓	✓	✓
	Refined	✓	✓	✓	✓	✓	-	_	_			✓	✓	✓	✓		-	-	-	-
		V	Fp	Fa	P	T	Te	A				P-A	P-Fp	P-T	P-V	A-Fa	T-Te			
DBLP	Vanilla	√.	√.	✓	-	-	✓	-				√.	-	√.	√.	-	-			
	Refined	√	√_	-	-	-	-	-				√	✓	✓	✓	-	-			
		T	A	S	P							P-P	P-A	P-S	P-T					
ACM	Vanilla	√.	-	-	-							√.	✓.	√.	√.					
	Refined	√	-	_	_							√	<u>√</u>	√	✓					
		F	E	Y	P	A	I					P-P	P-A	P-F	A-I	<u>P-E</u>	P-Y			
OGBN-MAG		✓,	√,	-	-	✓	√					√	√,	√,	√,	-	-			
	Refined	✓	√	√	-	_	_					✓	✓	✓	✓	√	-			
		I	P	В	R	C						I-I _v	$I-I_p$	I-R	I-P	P-C	P-B			
Amazon	Vanilla	-	✓	✓	✓	✓						✓	✓	-	-	-	-			
-	Refined	✓	-	_	-	-						✓	✓	✓	✓	-	-			
		T	A	U								A-T	U-U	U-A						
LastFM	Vanilla	✓	✓	✓								✓	✓	✓						
	Refined	✓	-	-								✓	-	-						
		C	G	$\mathbf{E_d}$	D	\mathbf{S}	$\mathbf{E}_{\mathbf{s}}$	$\mathbf{E}_{\mathbf{g}}$	$\mathbf{E}_{\mathbf{c}}$			C-S	C-C	G-S	G-G		C-D	$D-E_d$	C-G	$S-E_s$
PubMed	Vanilla	✓	✓	-	√	✓	-	-	-			✓	✓	✓	✓		✓	-	✓	-
	Refined	✓	✓	✓	✓	✓	_	_	_			-		-	-		_	-	-	-

4.3 How the Atomic View Benefits HGNNs? (RQ2)

Table 2 summarizes the refined schemas from *atomic HINs* using sRGCN across datasets derived from the proposed GA algorithm, comparing with the *vanilla* schemas. Details on how each vanilla schema corresponds to specific selections are provided in Appendix A.1. We make the following observations:

- Obs 1: Attribute atomization introduces meaningful relations. Refined schemas frequently
 preserve relations created when attributes are converted into entities—relations absent in vanilla
 schemas but induced through atomization. This pattern consistently appears in IMDB, DBLP,
 OGBN-MAG, and Amazon, where refined schemas select edges from atomized attributes, indicating that such induced relations provide valuable semantic signals for HGNNs.
- Obs 2: Entities can act as strong attributes. In Amazon, for example, *item* is originally described by attributes such as *price* or *sales-rank*. Refined schemas, however, benefit from directly learning ID embeddings for *item* nodes, while still retaining attributes like *price* and *sales-rank*. Similar effects are also observed in IMDB and PubMed.
- Obs 3: Relations from numerical attributes can be surprisingly useful. We atomize all attributes, including numerical ones, which induce relations corresponding to dense adjacencies. Although such dense edges may appear unintuitive, many are consistently selected in refined schemas (highlighted with double underlines in Table 2). A plausible explanation is that these edges encode similarity relations through metapaths (Sun et al., 2011). For example, the

Table 3: Performance of HGNNs under different schema variants. **Vanilla**: the original benchmark schema. **Refined(HGNN)**: schema refined using the corresponding HGNN. The best results within each HGNN are in **bold**, the second-best are <u>underlined</u>, and global best across all HGNNs are highlighted with <u>blue bold</u>.

Schema		IM	DB	Free	base	Ama	zon	OGBN-MAG	
HGNN		Macro-F1	Micro-F1	Macro-F1	Micro-F1	ROC-AUC	MRR	Acc. (Val)	Acc. (Test)
sRGCN	Vanilla	67.64±0.41	70.05 ± 0.50	52.13±1.78	67.09±0.43	95.94 ± 0.28	98.43 ± 0.18	56.73±0.21	54.63 ± 0.23
SKUCIN	Refined (sRGCN)	68.97 ± 0.09	71.20 ± 0.17	55.40±1.25	67.32 ± 0.66	97.85 ± 0.07	99.26±0.05	57.35±0.12	55.21 ± 0.23
	Vanilla	63.53±1.36	67.36±0.57	47.72 ± 1.48	66.29 ± 0.45	93.40 ± 0.62	96.94±0.29	OOM	OOM
SimpleHGN	Refined (sRGCN)	65.89 ± 0.67	68.60 ± 1.13	53.45 ± 1.88	67.83 ± 0.33	96.50 ± 0.87	98.61 ± 0.49	OOM	OOM
	Refined (SimpleHGN)	$\overline{67.38 \pm 0.80}$	70.02 ± 0.62	53.51±1.39	67.94 ± 0.80	97.40 ± 1.11	99.05 ± 0.08	OOM	OOM
	Vanilla	67.10±0.60	69.79 ± 0.52	40.01 ± 8.26	62.70 ± 0.77	94.12±0.58	97.93 ± 0.46	56.61±0.11	54.53 ± 0.20
PSHGCN	Refined (sRGCN)	67.89 ± 0.56	$69.87{\scriptstyle\pm1.04}$	45.53 ± 3.02	65.66 ± 0.25	96.73 ± 0.53	98.79 ± 0.23	57.65 ± 0.18	55.28 ± 0.19
	Refined (PSHGCN)	$67.89{\scriptstyle\pm0.56}$	$69.87{\scriptstyle\pm1.04}$	48.19 ± 1.45	66.36 ± 0.75	97.13 ± 0.11	98.91 ± 0.05	57.65 ± 0.09	55.34 ± 0.21

paper-author-paper metapath in the citation network captures co-authorship, while the paper-embedding-paper metapath in OGBN-MAG approximates paper similarity, akin to dot-product signals between embeddings.

4.4 IS SCHEMA REFINEMENT NECESSARY? (RQ3)

Having examined the role of *atomic HINs* and attribute atomization, we now turn to the necessity of *schema refinement*. Using the refined schemas reported in Table 2, we analyze how selectively retaining or discarding node and edge types affects performance.

- Obs 4: Schema refinement remains important even for fully atomic schemas. Some datasets, such as Freebase and LastFM, are already in atomic form and therefore cannot benefit from further atomization. Nevertheless, simple refinement strategies—selectively dropping node or edge types—still yield significant improvements, highlighting the value of schema selection beyond atomization alone.
- Obs 5: Schema refinement for link prediction favors pruning relations, even target ones. For link prediction tasks, refinement often requires more aggressive relation removal. In LastFM, dropping the user–artist relation (the prediction target) improves results, while in PubMed, discarding all edges yields the best performance. This behavior aligns with the over-smoothing effect observed by Butler et al. (2023): link prediction is more sensitive to excessive connectivity, whereas node classification typically benefits from homophily. Pruning relations that reduce connectivity can therefore lead directly to performance gains.

4.5 GENERALIZATION OF REFINED SCHEMAS (RQ4)

We now examine the quantitative benefits of schema refinement and its transferability across different HGNNs (Table 3). Beyond sRGCN, we evaluate two additional models: SimpleHGN, a representative attention-based HGNN, and PSHGCN, a recent precomputation-based HGNN that models SHGC without first-order approximation. Further discussion of these models and their connections to SHGC is provided in Appendix C.

- Obs 6: Schema variants yield substantial performance differences. Across all three HGNNs, refined schemas consistently outperform vanilla ones (first row of each HGNN), showing that schema choice can be as influential as model architecture.
- Obs 7: Refined schemas transfer effectively across HGNNs. Schemas discovered with sRGCN generalize well to other HGNNs (second row of each HGNN), often delivering strong results without re-optimization. A notable case is PSHGCN on OGBN-MAG, where the transferred schema achieves even better performance.
- Obs 8: Subsequent refinement provides small but consistent gains. Re-optimizing schemas for each HGNN (third row of each HGNN) yields additional improvements, though the margin is modest compared to the leap from vanilla to refined (sRGCN). This suggests that schemas found with sRGCN are already near-optimal for other HGNNs.

4.6 HOW EFFICIENT IS SCHEMA REFINEMENT AS A SEARCH PROBLEM? (RQ5)

We assess the efficiency of schema refinement by tracking performance over search trials (Figure 3). Starting from the vanilla schema, refinement converges rapidly even on relation-rich datasets. For

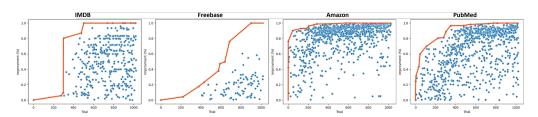


Figure 3: Performance improvement across schema search trials on four datasets.

instance, IMDB and PubMed reach near-optimal performance within 512 trials, despite the exhaustive search space containing up to 2^{19} and 2^{22} candidates, respectively.

5 RELATED WORK

5.1 HETEROGENEOUS GRAPH NEURAL NETWORKS

Heterogeneous graph neural networks (HGNNs) extend GNNs to exploit the semantic richness of heterogeneous information networks (HINs). Early models such as HAN (Wang et al., 2019) and MAGNN (Fu et al., 2020) rely on hand-crafted metapaths to define composite relations. RGCN (Schlichtkrull et al., 2018), HGT (Hu et al., 2020b), and Simple-HGN (Lv et al., 2021) instead directly model heterogeneity through relation-specific transformations or attention. GTN (Yun et al., 2019), MHGCN (Yu et al., 2022), and RE-GNN (Wang et al., 2023) further learn soft weights over edge types, effectively performing differentiable subgraph or metapath selection. More recently, SeHGNN (Yang et al., 2023), PSHGCN (He et al., 2024), and LMSPS (Li et al., 2024) precompute propagation along metapaths, enabling mini-batch training for large-scale HINs. Together, these methods capture relational information at the level of relations, metapaths, or soft subgraphs, optimized through manual design or differentiable learning.

5.2 HETEROGENEOUS GRAPH SCHEMA DESIGN

Compared to HGNN architectures, relatively little attention has been devoted to the construction of heterogeneous graph schemas. Common strategies include augmenting the schema with metapath-based edges (Wang et al., 2019; Fu et al., 2020; Hu et al., 2024) or introducing metapath-derived features (Lv et al., 2021; Fey et al., 2024). Benchmark datasets implement these strategies in different ways. For example, HGB (Lv et al., 2021) typically assigns identity matrices or metapath-based features to attribute-less entities, and selectively promotes some attributes (e.g., actors, keywords) while leaving others (e.g., language, country) as plain features. OGB (Hu et al., 2020a) similarly averages embeddings of textual terms into paper nodes but omits constructing the corresponding term nodes and relations. RelBench (Fey et al., 2024) moves toward more systematic construction, yet still admits multiple valid schemas. In contrast, we provide a **principled framework for schema design**. The atomic HIN unifies ad-hoc practices into a canonical representation where all design choices are explicit. On this foundation, schema refinement becomes a systematic, optimizable process rather than a heuristic one—establishing not only a complement to HGNN architectures but also a basis for rigorous benchmarking and schema-aware model development.

6 CONCLUSION

We introduced the *atomic HIN*, grounded in the principle of *entity-attribute duality*, where attributes can be atomized as entities and entities can in turn serve as attributes of others. This duality maximizes expressiveness by making schema choices explicit. On this foundation, we proposed schema refinement, a systematic procedure for selecting or discarding node- and edge-types to yield task-specific schemas. Across eight datasets, even a simplified RGCN (sRGCN) trained on refined *atomic HINs* achieves state-of-the-art performance, with further gains from advanced HGNNs. These results establish *atomic HINs* as a powerful representation for heterogeneous graph learning and highlight schema as a central dimension of HGNN design and evaluation. Looking ahead, we hope this work provides a foundation for principled schema-aware learning, automated schema discovery, and next-generation HGNN architectures.

REPRODUCIBILITY

To ensure reproducibility and facilitate future research, we provide detailed descriptions of the experimental setup in Appendix A.3 and the hyperparameter search space in Appendix A.4. The complete source code, together with processed datasets and search configurations, is temporarily hosted at https://anonymous.4open.science/r/AtomHIN-0D19. Datasets must be downloaded separately from the supplementary materials. We commit to releasing it publicly upon acceptance to enable transparent benchmarking and follow-up work.

REFERENCES

- HetRec '11: Proceedings of the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450310277.
- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2623–2631, 2019.
- Landon Butler, Alejandro Parada-Mayorga, and Alejandro Ribeiro. Convolutional learning on multigraphs. *Trans. Sig. Proc.*, 71:933–946, January 2023. ISSN 1053-587X. doi: 10.1109/TSP.2023. 3259144. URL https://doi.org/10.1109/TSP.2023.3259144.
- Yukuo Cen, Xu Zou, Jianwei Zhang, Hongxia Yang, Jingren Zhou, and Jie Tang. Representation learning for attributed multiplex heterogeneous network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1358–1368. ACM, 2019.
- Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. Measuring and relieving the oversmoothing problem for graph neural networks from the topological view. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3438–3445, Apr. 2020. doi: 10.1609/aaai.v34i04. 5747. URL https://ojs.aaai.org/index.php/AAAI/article/view/5747.
- K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi: 10.1109/4235.996017.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, 2016. URL https://arxiv.org/abs/1606.09375.
- Matthias Fey, Weihua Hu, Kexin Huang, Jan Eric Lenssen, Rishabh Ranjan, Joshua Robinson, Rex Ying, Jiaxuan You, and Jure Leskovec. Position: Relational deep learning graph representation learning on relational databases. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 13592–13607. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/fey24a.html.
- Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proceedings of The Web Conference 2020*, WWW '20, pp. 2331–2341, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450370233. doi: 10.1145/3366423.3380297. URL https://doi.org/10.1145/3366423.3380297.
- Mingguo He, Zhewei Wei, Shikun Feng, Zhengjie Huang, Weibin Li, Yu Sun, and Dianhai Yu. Spectral heterogeneous graph convolutions via positive noncommutative polynomials. In *Proceedings of the ACM Web Conference 2024*, WWW '24, pp. 685–696, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400701719. doi: 10.1145/3589334.3645515. URL https://doi.org/10.1145/3589334.3645515.

Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, YongDong Zhang, and Meng Wang. Lightgen: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '20, pp. 639–648, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380164. doi: 10.1145/3397271.3401063. URL https://doi.org/10.1145/3397271.3401063.

- Jun Hu, Bryan Hooi, and Bingsheng He. Efficient heterogeneous graph learning via random projection. *IEEE Transactions on Knowledge and Data Engineering*, pp. 1–14, 2024. doi: 10.1109/TKDE.2024.3434956.
- Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: datasets for machine learning on graphs. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA, 2020a. Curran Associates Inc. ISBN 9781713829546.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. Heterogeneous graph transformer. In *Proceedings of The Web Conference 2020*, WWW '20, pp. 2704–2710, New York, NY, USA, 2020b. Association for Computing Machinery. ISBN 9781450370233. doi: 10.1145/3366423. 3380027. URL https://doi.org/10.1145/3366423.3380027.
- Sourabh Katoch, Sumit Singh Chauhan, and Vijay Kumar. A review on genetic algorithm: past, present, and future. *Multimedia Tools and Applications*, 80(5):8091–8126, 2021. ISSN 1573-7721. doi: 10.1007/s11042-020-10139-6. URL https://doi.org/10.1007/s11042-020-10139-6.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Chao Li, Zijie Guo, Qiuting He, Hao Xu, and Kun He. Long-range meta-path search on large-scale heterogeneous graphs, 2024.
- Qingsong Lv, Ming Ding, Qiang Liu, Yuxiang Chen, Wenzheng Feng, Siming He, Chang Zhou, Jianguo Jiang, Yuxiao Dong, and Jie Tang. Are we really making much progress? revisiting, benchmarking and refining heterogeneous graph neural networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, KDD '21, pp. 1150–1160, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383325. doi: 10.1145/3447548.3467350. URL https://doi.org/10.1145/3447548.3467350.
- Qiheng Mao, Zemin Liu, Chenghao Liu, and Jianling Sun. Hinormer: Representation learning on heterogeneous information networks with graph transformer. In *Proceedings of the ACM Web Conference 2023*, WWW '23, pp. 599–610, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394161. doi: 10.1145/3543507.3583493. URL https://doi.org/10.1145/3543507.3583493.
- Yoshihiko Ozaki, Yuki Tanigaki, Shuhei Watanabe, and Masaki Onishi. Multiobjective treestructured parzen estimator for computationally expensive optimization problems. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, GECCO '20, pp. 533–541, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450371285. doi: 10.1145/3377930.3389817. URL https://doi.org/10.1145/3377930.3389817.
- Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam (eds.), *The Semantic Web*, pp. 593–607, Cham, 2018. Springer International Publishing. ISBN 978-3-319-93417-4.
- Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu, and Tianyi Wu. Pathsim: meta path-based top-k similarity search in heterogeneous information networks. *Proc. VLDB Endow.*, 4(11):992–1003, August 2011. ISSN 2150-8097. doi: 10.14778/3402707.3402736. URL https://doi.org/10.14778/3402707.3402736.

- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJXMpikCZ. accepted as poster.
- Junfu Wang, Yuanfang Guo, Liang Yang, and Yunhong Wang. Enabling homogeneous gnns to handle heterogeneous graphs via relation embedding. *IEEE Transactions on Big Data*, 9:1697–1710, 2023.
- Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The World Wide Web Conference*, WWW '19, pp. 2022–2032, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450366748. doi: 10. 1145/3308558.3313562. URL https://doi.org/10.1145/3308558.3313562.
- Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 6861–6871. PMLR, 2019.
- Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In *International Conference on Learning Representations*, 2014. URL https://api.semanticscholar.org/CorpusID: 2768038.
- Carl Yang, Yuxin Xiao, Yu Zhang, Yizhou Sun, and Jiawei Han. Heterogeneous network representation learning: A unified framework with survey and benchmark. *TKDE*, 2020.
- Xiaocheng Yang, Mingyu Yan, Shirui Pan, Xiaochun Ye, and Dongrui Fan. Simple and efficient heterogeneous graph neural network. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(9):10816–10824, Jun. 2023. doi: 10.1609/aaai.v37i9.26283. URL https://ojs.aaai.org/index.php/AAAI/article/view/26283.
- Pengyang Yu, Chaofan Fu, Yanwei Yu, Chao Huang, Zhongying Zhao, and Junyu Dong. Multiplex heterogeneous graph convolutional network. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2377–2387, 2022.
- Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph transformer networks. In *Advances in Neural Information Processing Systems*, pp. 11960–11970, 2019.
- Ziang Zhou, Jieming Shi, Renchi Yang, Yuanhang Zou, and Qing Li. Slotgat: slot-based message passing for heterogeneous graphs. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.

A EXPERIMENTAL SETUP

A.1 DATASETS

We provide detailed descriptions of the eight heterogeneous datasets used in our experiments. For each dataset, we highlight the vanilla schema, the preprocessing conventions used in prior benchmarks, and how our atomization modifies the schema.

• IMDB (Lv et al., 2021): A movie network with entities including movie, actor, director, keyword, word, color, country, content rating, and language. In the vanilla setting, only movie, actor, director, and keyword are retained as explicit node types. The remaining entities—word, color, country, content rating, and language—are collapsed into attributes of the primary nodes (movie, actor, or director), while keyword nodes are assigned unique identity vectors. Within our schema refinement framework, this corresponds to selecting keyword, word, color, country, content rating, and language, while the entities movie, actor, and director are unselected. Note that keyword and word are distinct: for example, a keyword might be "coming of age," while the corresponding words are "coming," "of," and "age," with frequency counts. For clarity in the main text, we consistently refer to keywords rather than words to avoid terminological confusion, though both choices carry practical performance impact in downstream tasks. In the atomic form, all attributes

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692 693

694

696

697

698

699

700

701

Table 4: Dataset statistics for benchmark and atomic HIN.

#Edge-Types

#Nodes

24.909

30,743

10 942

1,939,879

#Node-Types

Atomic HIN

#Edges

94.384

1,057,688

263,623

279 221

21,740,578

#Edge-Types Dense Adj.

36

4.932×16

7,723×50

736,389×128

Benchmark (vanilla)

43.321

1,057,688

119,783

21,111,007

#Node-Types #Edges

Target (type / etype) #Classes #Nodes Node Classification (NC) IMDB 21.420 DBLP Author 26,128 ACM 10 942 OGBN-MAG Paper 349 1,939,743 Link Prediction (LP) 10,099 Amazon item-item

Amazon item-item — 10,099 1 121,470 2 11,256 5 151,729 6 $10,099\times 1$ LastFM artist-user — 20,612 3 111,796 3 20,612 3 111,796 3 — - PubMed disease-disease — 63,109 4 233,047 10 63,909 8 233,047 14 $200\times 13,561$ are expanded into explicit node types, ensuring that every entity—including word, color, country, $content\ rating$, and language—is represented relationally.

- Freebase (Lv et al., 2021): A knowledge graph with 8 node types and 36 edge types, constructed without attributes. Since no attributes are available, the vanilla and atomic schemas coincide.
- **DBLP** (Lv et al., 2021): A bibliographic network with node types *author*, *paper*, *venue*, and *term*. In the vanilla form, authors and papers are assigned undocumented multi-hot sparse features; venues have no attributes, while terms are initialized with 50-dim PLM embeddings. Thus we atomize them into dedicated *author-feat* and *paper-feat* nodes, together with *term-feat* nodes for pretrained embeddings.
- ACM (Lv et al., 2021): A citation network with node types *paper*, *subject*, *author*, and *term*. In the vanilla form, only terms are directly assigned unique identity vectors; other node types (*subject*, *author*, *paper*) are preprocessed with propagated term IDs as their attributes. Thus, only *term* nodes are selected in the vanilla schema. The benchmark schema is already atomic, and no further atomization is allowed.
- OGBN-MAG (Hu et al., 2020a): A large-scale academic graph targeting venue classification. Papers are annotated with publication year and pretrained text embeddings. In the vanilla form, embeddings are obtained by mean-pooling over term nodes that are not explicitly constructed. Following current practice (e.g., Yang et al. (2023)), we assign each node type a 256-dim random embedding table as an approximation for larget-scale node types (paper, author, institution, field-of-study). In addition, we atomize the 128-dim PLM embeddings into feature nodes and expand year information into discrete year nodes, restricted to training papers due to the chronological split.
- Amazon (Cen et al., 2019): An e-commerce graph with a single node type *item* and edge types *co-view* and *co-purchase*. Items are annotated with attributes including price, sales-rank, category, and brand. The *price* attribute is a one-dimensional numerical feature. Direct normalization produces trivial results (all values collapse to ± 1); to avoid this, we append a dummy column with values equal to the mean absolute price, ensuring non-trivial normalization. All attributes are then atomized into corresponding feature nodes.
- LastFM (Het, 2011): A music network with node types *user*, *artist*, and *tag*, targeting user-artist edge prediction. No attributes are provided in the benchmark form, so the graph is already atomic.
- **PubMed** (Yang et al., 2020): A biomedical knowledge graph with node types *disease*, *gene*, *chemical*, and *species*. All nodes are provided with 256-dim pretrained embeddings, which are often replaced by learnable ID embeddings in literatures (we thus mark this variant as *vanilla* in Table 2 as most baselines adopt such variant). We atomize these embeddings into separate feature nodes for each entity type, yielding a fully relational schema.

A.1.1 ALIASES FOR NODE TYPES

Aliases for node types used in Table 2:

- IMDB: N = numerical, R = content rating, D = director, W = word, C = country, L = language, A = actor, O = color, K = keyword.
- Freebase: Bk = book, Or = organization, Bs = business, Mu = music, Sp = sports, Lo = location, Fi = film, Pe = people.
- **DBLP:** A = author, P = paper, V = venue, T = term, A_f = author-feature, P_f = paper-feature, T_f = term-feature.

- ACM: P = paper, A = author, S = subject, T = term.
- Amazon: I = item, P = price, B = brand, R = sales-rank, C = category.
- LastFM: U = user, A = artist, T = tag.
- **PubMed:** D = disease, G = gene, C = chemical, S = species, D_f = disease-feature, G_f = gene-feature, C_f = chemical-feature, S_f = species-feature.
- **OGBN-MAG:** P = paper, A = author, I = institution, F = field-of-study, E = term-embedding, Y = year.

A.2 BASELINES HGNNS

We briefly describe the HGNN baselines used in our experiments:

- **RGCN** (Schlichtkrull et al., 2018): One of the earliest HGNNs, extending GCNs to relational graphs by introducing relation-specific feature transformations. Each edge type is assigned a distinct projection matrix, enabling type-aware message passing.
- **HGT** (Hu et al., 2020b): An early attention-based HGNN. It models heterogeneity through relation-specific attention with node-type-aware key/query projections, combined with learnable relation-specific priors to capture the varying importance of different relations.
- **SimpleHGN** (Lv et al., 2021): A simple extension of GAT without heavily modeling heterogeneity. It incorporates relation-aware information through relation embeddings, achieving competitive performance.
- **HINormer** (Mao et al., 2023): Adapts Graph Transformers to HINs by combining them with GCN-style propagation. It first encodes local structure through neighborhood aggregation, then applies a heterogeneous relation encoder to model relation-specific information.
- **SeHGNN** (Yang et al., 2023): Extends SGC-style precomputation to HINs, enabling efficient precomputation-based training. This design allows mini-batch learning without information loss from subgraph sampling or GPU memory overhead, making it well-suited for large-scale settings.
- **SlotGAT** (Zhou et al., 2023): A GAT-based model for heterogeneous graphs. It conducts message passing separately across node-type-specific "slots," preserving distinct semantics in different feature spaces.
- **PSHGCN** (He et al., 2024): A spectral HGNN that models heterogeneous graphs under positive semi-definite constraints on spectral filters. Its framework also extends to precomputation-based settings, enabling scalable and efficient training on large-scale HINs.

A.3 EVALUATION SETTING

For node classification datasets (IMDB, Freebase, DBLP, ACM, OGBN-MAG), we follow the official splits and evaluation metrics provided by each benchmark. On OGBN-MAG, we report classification accuracy on both validation and test sets, following the official protocol, and present the mean and standard deviation over 10 runs. On the other datasets, we report both Macro-F1 and Micro-F1, averaged over 5 runs.

For link prediction datasets (Amazon, LastFM, PubMed), the task is cast as binary classification on node pairs. We evaluate with ROC-AUC and MRR, using negative samples drawn from 2-hop neighbors of each positive pair, following the protocol of Lv et al. (2021). All datasets use official splits except PubMed, which was originally introduced by Yang et al. (2020). We found its official split suffers from severe distribution shift, and thus re-split the dataset (with the same ratio). All experiments are conducted on a single RTX A6000 GPU with 48GB memory or smaller GPUs.

A.4 HYPERPARAMETER SETTING

As discussed in Section 3.2, we treat schema parameters and model depth as hyperparameters. Hyperparameter optimization is carried out using optuna (Akiba et al., 2019). We adopt a two-stage tuning process: (1) schema tuning with a budget of 1024 trials, followed by (2) HGNN hyperparameter tuning with 256 trials.

Schema Tuning. We use a genetic algorithm (Deb et al., 2002) with default parameters, except that mutation probability is set to 1/8 for datasets with more than eight schema parameters. The search is initialized from the vanilla schema configuration, and explores 1024 candidate points.

HGNN Hyperparameter Tuning. We apply the Tree-structured Parzen Estimator (TPE) (Ozaki et al., 2020) implemented in optuna, using default parameters. For each model, we define the following search spaces:

- **sRGCN:** number of layers, hidden dimension, number of heads, dropout rate, edge dropout (or input dropout for precomputation-based setting on OGBN-MAG), learning rate, weight decay, decoder MLP layers.
- SimpleHGN (Lv et al., 2021): number of layers, hidden dimension, number of heads, dropout rate, edge dropout, learning rate, weight decay.
- PSHGCN (He et al., 2024): number of layers/hops, hidden dimension, embedding dimension, dropout rate, input dropout, learning rate, weight decay.

B DETAILS OF PRE-PROPAGATION FEATURE INITIALIZATION

This section provides further motivation and proofs for the lemmas stated in Section 3.3, along with additional technical details.

B.1 Proof of Lemma 3.1

Lemma (Independence of Selections). With pre-propagation, for any node type τ_i and any selected type τ_j with $\beta_{\tau_j} = 1$, $\hat{\mathbf{I}}_{\tau_i} \bar{\mathbf{Z}} \hat{\mathbf{I}}_{\tau_j} \neq \mathbf{0}$. Hence, node-type selection is independent of edge-type selection: even if all edges incident to τ_j are removed, no dependency arises.

Proof. To prove independence, we must show that whenever $\beta_{\tau_j}=1$, there must exist at least one non-zero term in \bar{Z} that connects τ_j to every other node type τ_i . Formally, if $\beta_{\tau_j}=1$,

$$\hat{\mathbf{I}}_{\tau_i} \, \bar{\mathbf{Z}} \, \hat{\mathbf{I}}_{\tau_i} \neq \mathbf{0}, \qquad \forall \tau_i \in \{1, \dots, |\mathcal{T}|\},$$
 (5)

where \bar{Z} is the SHGC polynomial (from Equation 2).

Expanding the polynomial gives

$$\bar{\mathbf{Z}} = \mathbf{H}(\mathbf{I}, \alpha_{1} \mathbf{S}_{1}, \dots, \alpha_{R} \mathbf{S}_{R}; \Theta) \mathbf{X} \left(\beta_{1} \hat{\mathbf{I}}_{1} + \dots + \beta_{|\mathcal{T}|} \hat{\mathbf{I}}_{|\mathcal{T}|}\right)$$

$$= \left(\theta_{0} \mathbf{I} + \sum_{\ell=1}^{L} \sum_{r_{1}, \dots, r_{\ell}} \theta_{r_{1}, \dots, r_{\ell}} (\alpha_{r_{1}} \cdots \alpha_{r_{\ell}}) (\mathbf{S}_{r_{1}} \cdots \mathbf{S}_{r_{\ell}})\right) \left(\mathbf{I} + \sum_{\tau_{i} \neq \tau_{j}} \tilde{\mathbf{A}}_{\langle \tau_{i}, *, \tau_{j} \rangle}\right) (\beta_{1} \hat{\mathbf{I}}_{1} + \dots + \beta_{|\mathcal{T}|} \hat{\mathbf{I}}_{|\mathcal{T}|}).$$

The second factor in Equation 6 arises from pre-propagation and includes terms of the form

$$\hat{\boldsymbol{I}}_{\tau_i} \Big(\boldsymbol{I} + \sum_{\tau_i \neq \tau_j} \tilde{\boldsymbol{A}}_{\langle \tau_i, *, \tau_j \rangle} \Big) \hat{\boldsymbol{I}}_{\tau_j} \neq 0, \quad \forall \, \tau_i, \tau_j \in \{ 1, \dots, |\mathcal{T}| \}.$$

These terms remain non-zero regardless of the values of α_r .

Hence, even if all relations incident to τ_j are deselected (i.e., $\alpha_r = 0$), the pre-propagated features still allow τ_j to contribute to other node types. This shows that node-type and edge-type selections are independent once pre-propagation is applied.

B.2 Proof of Lemma 3.2

Lemma (Neutrality of Pre-propagation). Consider an SHGC with row-normalized adjacencies $S_r = \tilde{A}_r$ as shift operators and convolution order L sufficiently large. If raw identity features $X_0(1,\ldots,1) = I$ are replaced with pre-propagated features $X(1,\ldots,1)$, the effect is merely a reparameterization of the filter coefficients $\{\theta_{r_1,\ldots,r_\ell}\}_{r_1,\ldots,r_\ell}$. Hence, pre-propagation does not alter the expressive power of the model.

Proof. Let $S_r = \tilde{A}_r$ denote row-normalized adjacency operators. For sufficiently large convolution depth L, the SHGC representation on identity features I can be written as

$$ar{oldsymbol{Z}} = heta_0 oldsymbol{I} + \sum_{\ell=1}^L \sum_{r_1,...,r_\ell} heta_{r_1,...,r_\ell} \left(oldsymbol{S}_{r_1} \cdots oldsymbol{S}_{r_\ell}
ight) oldsymbol{I}.$$

When pre-propagation is applied, the identity features are replaced by

$$\boldsymbol{X}(1,\ldots,1) = \left(\boldsymbol{I} + \sum_{\tau_i \neq \tau_j} \tilde{\boldsymbol{A}}_{\langle \tau_i, *, \tau_j \rangle} \right) \boldsymbol{I},$$

so that

$$\bar{Z}' = \left(\theta_0 \mathbf{I} + \sum_{\ell=1}^{L} \sum_{r_1, \dots, r_\ell} \theta_{r_1, \dots, r_\ell} \left(\mathbf{S}_{r_1} \cdots \mathbf{S}_{r_\ell} \right) \right) \mathbf{X}(1, \dots, 1)$$

$$= \left(\theta_0 \mathbf{I} + \sum_{\ell=1}^{L} \sum_{r_1, \dots, r_\ell} \theta'_{r_1, \dots, r_\ell} \left(\mathbf{S}_{r_1} \cdots \mathbf{S}_{r_\ell} \right) \right) \mathbf{I}, \tag{7}$$

where the coefficients $\theta'_{r_1,\ldots,r_\ell}$ are obtained by reparametrizing $\theta_{r_1,\ldots,r_\ell}$ after expanding the product with $X(1,\ldots,1)$.

Equation equation 7 shows that pre-propagation merely reweights the convolution coefficients, but the span of operators $\{S_{r_1}\cdots S_{r_\ell}\}_{r_1,\dots,r_\ell}$ remains unchanged. Thus, the filter space of SHGC with pre-propagated features coincides with that of SHGC on raw identity features, provided L is sufficiently large.

Therefore, pre-propagation is neutral with respect to expressive capacity, while simultaneously enabling the independence property established in Lemma 3.1.

B.3 REMARKS

We highlight two further observations related to the role of pre-propagation:

- The proposed pre-propagation is not the only possible way to resolve the dependency between node-type and edge-type selection. For example, in node-level tasks with a designated target type τ_t , it suffices to ensure $\hat{I}_{\tau_t}\bar{Z}\hat{I}_{\tau_j} \neq 0$ for the selected type τ_j . Our formulation, however, provides a general and convenient solution that applies uniformly across tasks.
- Pre-propagation does not introduce additional learnable parameters. Feature transformation operates on $\boldsymbol{X}(\beta_1,\ldots,\beta_T)\boldsymbol{W}$, and the effective number of parameters in \boldsymbol{W} depends only on the non-zero columns of $\boldsymbol{X}(\cdot)$, i.e., $\sum_{\tau=1}^{|\mathcal{T}|}\beta_\tau|\mathcal{V}_\tau|$. Thus, the number of embeddings is determined entirely by node-type selection, independent of whether pre-propagation is applied.

B.4 GENERALIZED UNION AGGREGATION.

We first demonstrate pre-propagation using row-normalized adjacencies \tilde{A}_r , which corresponds to mean or weighted-sum pooling:

$$\bar{\boldsymbol{x}}_i = \sum_{j \in \mathcal{N}_i^{(r)}} \tilde{A}_r[i, j] \, \boldsymbol{x}_j, \tag{8}$$

where $\mathcal{N}_i^{(r)}$ denotes the neighbors of node i under relation r.

In practice, however, union-based aggregation is often adopted for initializing features, as in Lv et al. (2021). In the unweighted case, this is equivalent to an element-wise maximum across neighbors:

$$\bar{x}_i[k] = \max_{j \in \mathcal{N}_i^{(r)}} x_j[k]. \tag{9}$$

We extend this to a generalized union operator that performs element-wise maximization while ignoring sign. This variant naturally handles weighted adjacencies and empirically yields more

Algorithm 1 Pre-propagation of features

864

892

893

894 895 896

897

899 900

901

902

903 904

905

906

907

908

909

910

911 912

913914915

916

917

```
865
                 Input: \{\tilde{A}_r\}_r: row-normalized adjacency matrices of relations in the heterogeneous graph
866
                 Input: \{X_{\tau}\}_{\tau=1}^{|\mathcal{T}|}: raw feature matrices for each node type
867
                Output: \{X_{\tau}'\}_{\tau=1}^{|\mathcal{T}|}: pre-propagated feature matrices
868
                   1: procedure PREPROPAGATE(\{\tilde{A}_r\}_r, \{X_\tau\}_{\tau=1}^{|\mathcal{T}|})
870
                               Initialize \{ \boldsymbol{X}_{\tau}' \}_{\tau=1}^{|\mathcal{T}|} \leftarrow \{ \boldsymbol{0} \}_{\tau=1}^{|\mathcal{T}|} for \tau_j \leftarrow 1 to |\mathcal{T}| do
871
                  3:
872
                                      Construct \{X_{\tau_i}^{(\tau_j)}\}_{\tau_i=1}^{|\mathcal{T}|} where X_{\tau_j}^{(\tau_j)} \leftarrow X_{\tau_i} and X_{\tau_i}^{(\tau_j)} \leftarrow 0 for \tau_j \neq \tau_i
                  4:
873
                                      \{oldsymbol{Z}_{	au_i}\}_{	au_i=1}^{|\mathcal{T}|} \leftarrow 	ext{PropagateFeatures}(\{oldsymbol{	ilde{A}}_r\}_r, \{oldsymbol{X}_{	au_i}^{(	au_j)}\}_{	au_i=1}^{|\mathcal{T}|})
                  5:
874
                                      \begin{array}{c} \textbf{for } \tau_i \leftarrow 1 \textbf{ to } |\mathcal{T}| \textbf{ do } \\ X'_{\tau_i} \leftarrow X'_{\tau_i} + Z_{\tau_i} \\ \textbf{end for } \end{array}
                  6:
875
                  7:
876
                  8:
877
                               end for
                  9:
878
                               return \{X_{	au}'\}_{	au=1}^{|\mathcal{T}|}
                 10:
879
880
                               function PROPAGATEFEATURES(\{\tilde{A}_r\}_r, \{X_\tau\}_{\tau=1}^{|\mathcal{T}|})
                 11:
                                       while \exists 	au_i such that oldsymbol{X}_{	au_i} = oldsymbol{0} do
                 12:
                 13:
                                              for all A_r \in \{A_r\}_r do
883
                 14:
                                                     (\tau_s, \tau_t) \leftarrow source and target node types of relation r
884
                                                     if X_{	au_s} 
eq 0 and X_{	au_t} = 0 then
                 15:
885
                                                            X_{	au_t} \leftarrow \hat{A}_r X_{	au_s}
                 16:
886
                 17:
887
                 18:
                                              end for
                 19:
                                       end while
                                       return \{X_{\tau}\}_{\tau=1}^{|\mathcal{T}|}
                 20:
889
                 21:
                                end function
890
                 22: end procedure
891
```

stable performance when used for pre-propagation:

$$\bar{x}_i[k] = A_r[i, t] x_t[k], \qquad t = \arg \max_{j \in \mathcal{N}_i^{(r)}} A_r[i, j] x_j[k].$$
 (10)

Here the chosen neighbor t is the one with the maximum (magnitude) contribution on dimension k. Overall, we find that union-based pre-propagation tends to yield more consistent performance across datasets compared to mean or weighted-sum pooling with \tilde{A}_r .

B.5 IMPLEMENTATION.

Pre-propagation is performed prior to schema refinement. This procedure is independent of edgetype selection and thus only needs to be pre-computed once for each dataset.

The computation proceeds independently for each node type. For a given node type τ_j , we initialize its feature matrix as $\boldsymbol{X}_{\tau_j} = \hat{\boldsymbol{I}}_{\tau_j}$, i.e., identity features for nodes of type τ_j and zeros elsewhere. Propagation is then applied along any relation whose source node type already has non-empty features and whose destination node type is still empty. This process is repeated iteratively until all node types have received non-empty features, yielding the pre-propagated initialization $\boldsymbol{X}(1,\ldots,1)\hat{\boldsymbol{I}}_{\tau_j}$ for node type τ_j .

The full procedure is summarized in Algorithm 1.

C SHGC Perspective on HGNNs

This appendix provides detailed case analyses of how representative HGNNs relate to the Spectral Heterogeneous Graph Convolution (SHGC) formulation introduced in Equation 1. We show that several widely used models can be interpreted as first-order approximations of SHGC under different parameter-sharing schemes.

Table 5: Representative HGNNs interpreted under SHGC. Approximation indicates whether the model corresponds to a first-order approximation of SHGC or models SHGC directly. Shift operators: \tilde{A}_T denotes relation-specific row-normalization, and \hat{A}_T denotes cross-type row-normalization.

Model	Approximation	Shift Operator	Parameter Sharing
RGCN (Schlichtkrull et al., 2018) sRGCN (ours)	First-order	$egin{array}{c} ilde{A}_r \ ilde{A}_r \end{array}$	$egin{aligned} oldsymbol{W}_r^{(\ell)} &= \sum_b heta_{r,b} oldsymbol{W}_b^{(\ell)} \ oldsymbol{W}_r^{(\ell)} &= heta_r oldsymbol{I} \end{aligned}$
GTN (Yun et al., 2019)	First-order	$ ilde{m{A}}_r$	$oldsymbol{W}_r^{(\ell)} = heta_r^{(\ell)} oldsymbol{I}$
SimpleHGN (Lv et al., 2021) SeHGNN (Yang et al., 2023)	First-order	$egin{aligned} \hat{A}_r \ ilde{A}_r \end{aligned}$	$egin{aligned} oldsymbol{W}_r^{(\ell)} &= heta_r^{(\ell)} oldsymbol{W}^{(\ell)} \ oldsymbol{W}_{r_1,,r_\ell} \end{aligned}$
PSHGCN (He et al., 2024)	_	$ ilde{m{A}}_r$	$oldsymbol{W}_{r_1,,r_\ell} = heta_{r_1,,r_\ell} oldsymbol{W}$

Formally, extending SHGC (Equation 1) to multi-channel features $m{X} \in \mathbb{R}^{|\mathcal{V}| imes d}$ yields

$$Z = IXW_0 + \sum_{\ell=1}^{L} \sum_{r_1, \dots, r_{\ell}} S_{r_1} S_{r_2} \cdots S_{r_{\ell}} X W_{r_1, \dots, r_{\ell}},$$
(11)

where $W_{r_1,...,r_\ell} \in \mathbb{R}^{d \times d'}$ are generalized filter coefficients for multi-channel features and S_r denotes the shift operator associated with relation r.

C.1 First-Order Approximations

Relational GCN (RGCN). RGCN (Schlichtkrull et al., 2018) extends GCN (Kipf & Welling, 2017) to heterogeneous graphs. Following GCN, which adopts a first-order approximation of spectral convolution (i.e., L=1) and realizes higher-order convolution by stacking multiple layers, RGCN can likewise be viewed as a first-order approximation of SHGC:

$$m{H}^{(\ell+1)} = m{I}m{H}^{(\ell)}m{W}_0^{(\ell)} + \sum_{r=1}^{|\mathcal{R}|} ilde{m{A}}_rm{H}^{(\ell)}m{W}_r^{(\ell)},$$

where \tilde{A}_r denotes the normalized adjacency matrix for relation r.²

Even with this first-order approximation, the number of parameters $\{W_r^{(\ell)} \in \mathbb{R}^{d \times d'}\}_{r,\ell}$ can still be prohibitive in relation-rich graphs. To address this, RGCN introduces a basis-decomposition scheme for parameter sharing:

$$\boldsymbol{W}_r^{(\ell)} = \sum_{b=1}^B \theta_{r,b} \, \boldsymbol{W}_b,$$

where $B \in \mathbb{N}$ is the number of basis transformations, $W_b \in \mathbb{R}^{d \times d'}$ are shared across all relations, and $\theta_{r,b} \in \mathbb{R}$ are relation-specific coefficients.

Graph Transformer Network (GTN). GTN (Yun et al., 2019) learns weighted combinations of adjacency matrices to construct meta-relations, producing a meta-path adjacency A_P on which a GCN can be applied.

Formally, A_P is learned for an arbitrary meta-path of length at most L as

$$oldsymbol{A}_P = \prod_{\ell=1}^L \left(heta_0^{(\ell)} oldsymbol{I} + \sum_{r=1}^{|\mathcal{R}|} heta_r^{(\ell)} oldsymbol{A}_r
ight),$$

where self-loops are included to allow paths shorter than L hops. A GCN is then applied using the row-normalized adjacency \tilde{A}_P as the propagation operator:

$$Z = (I + \tilde{A}_P)XW.$$

²RGCN considers different normalization choices depending on the task: relation-specific normalization for node-level tasks, and cross-relation normalization for link-level tasks. Here, \tilde{A}_r corresponds to the relation-specific normalization.

Expanding \tilde{A}_P shows that Z can be written as

$$oldsymbol{Z} = heta_0' oldsymbol{I} oldsymbol{X} oldsymbol{W} + \sum_{\ell=1}^L \sum_{r_1, ..., r_\ell} heta_{r_1, ..., r_\ell}' ilde{oldsymbol{A}}_{r_1} \cdots ilde{oldsymbol{A}}_{r_\ell} oldsymbol{X} oldsymbol{W},$$

where each coefficient $\theta'_{r_1,\dots,r_\ell}$ is a product of up to L scalars $\{\theta_r^{(\ell)}\}_{r,\ell}$. For example, $\theta'_0=1+\prod_{\ell=1}^L\theta_0^{(\ell)}$ corresponds to the identity term.

This expansion demonstrates that GTN is another example of a first-order approximation of SHGC, but with stronger parameter sharing:

$$W_r^{(\ell)} = \theta_r^{(\ell)} I$$
.

In practice, GTN further introduces multiple sets of scalars $\theta_r^{(\ell,b)}$ for $b=1,\ldots,B$ to capture diverse meta-relational structures. We omit this extension here for clarity.

SimpleHGN. Attention-based HGNNs, such as SimpleHGN (Lv et al., 2021), extend GAT (Veličković et al., 2018) by introducing relation-aware attention through a relation-specific bias in the attention mechanism.

Empirically, attention models behave similarly to spectral graph convolution with row-normalization (Chen et al., 2020). Assuming further that the relation-specific bias acts effectively as a multiplicative scalar, SimpleHGN can be interpreted as a first-order SHGC approximation with parameter-sharing as:

$$\boldsymbol{W}_r^{(\ell)} = \theta_r^{(\ell)} \boldsymbol{W}^{(\ell)}, \tag{12}$$

where $W^{(\ell)} \in \mathbb{R}^{d \times d'}$ is a shared projection and $\theta_r^{(\ell)} \in \mathbb{R}$ is a relation-specific scalar.

Concretely, SimpleHGN incorporates edge-type information by introducing a learnable relation embedding $\boldsymbol{w}_r \in \mathbb{R}^{d_{\text{re}}}$ (with d_{re} a hyperparameter) and a projection $\boldsymbol{W}_{\text{rel}}$. For an edge (i,j) of type $r = \psi(\langle i,j \rangle)$, the modified attention score is

$$\hat{lpha}'_{i,j} = oldsymbol{a}^{ op} ig[oldsymbol{W}^{ op} oldsymbol{h}_i \parallel oldsymbol{W}^{ op} oldsymbol{h}_j \parallel oldsymbol{W}_{ ext{rel}}^{ op} oldsymbol{w}_r ig] \\ = \hat{lpha}_{i,j} + oldsymbol{a}_2^{ op} oldsymbol{W}_{ ext{rel}}^{ op} oldsymbol{w}_r,$$

where $\psi: \mathcal{V} \times \mathcal{V} \to \mathcal{R}$ is the edge-type indicator, $\hat{\alpha}_{i,j} = \boldsymbol{a}_1^{\top} [\boldsymbol{W}^{\top} \boldsymbol{h}_i \parallel \boldsymbol{W}^{\top} \boldsymbol{h}_j]$ is the vanilla GAT attention score, and $\boldsymbol{a} = \boldsymbol{a}_1 \parallel \boldsymbol{a}_2$. (For readability, we omit the layer index (ℓ) in this part; it is reintroduced in the update equations below.)

Although this formulation is an intuitive extension of GAT, the relation-specific term $a_2^\top W_{\text{rel}}^\top w_r$ can be reparametrized by a single scalar θ_r , yielding

$$\hat{\alpha}'_{i,j} = \hat{\alpha}_{i,j} + \theta_r.$$

Just as GTN introduces multiple scalar coefficients per relation, SimpleHGN extends to use multihead attention, each head learning its own $\theta_{r,b}$.

Following the empirical observation that attention resembles row-normalized spectral convolution (Chen et al., 2020), we may approximate the relational bias as multiplicative scalar edge weights $\hat{\theta}_r^{(\ell)}$. This yields the SimpleHGN update rule

$$\boldsymbol{h}_i^{(\ell+1)} = \sum_{r=1}^{|\mathcal{R}|} \sum_{j \in \mathcal{N}^{(r)}} \hat{\theta}_r^{(\ell)} \boldsymbol{W}^{(\ell) \top} \boldsymbol{h}_j^{(\ell)} \ + \ \boldsymbol{W}_0^{(\ell) \top} \boldsymbol{h}_i^{(\ell)},$$

where $\mathcal{N}_i^{(r)}$ denote the neighbors of *i*-th node via edge type r. Or equivalently in matrix form,

$$m{H}^{(\ell+1)} = \sum_{r=1}^{|\mathcal{R}|} \hat{ heta}_r^{(\ell)} \hat{m{A}}_r m{H}^{(\ell)} m{W}^{(\ell)} \ + \ I m{H}^{(\ell)} m{W}_0^{(\ell)},$$

where \hat{A}_r is the cross-type row-normalized adjacency matrix for relation r, approximating the soft-max normalization of GAT. This shows the parameter-sharing introduced earlier in Equation 12.

C.2 BEYOND FIRST-ORDER APPROXIMATIONS

While many HGNNs adopt first-order approximations, several recent works learn SHGC more directly without truncation. These models are often precomputation-based, computing $(S_{r_1}S_{r_2}\cdots S_{r_\ell})X$ offline. Such precomputation is advantageous when the number of relations is moderate, as it enables efficient mini-batch training: precomputed features $(S_{r_1}S_{r_2}\cdots S_{r_\ell})X$ can be sliced directly without subgraph sampling, thereby avoiding the trade-off between information loss (few hops or few sampled neighbors) and GPU memory overhead.

SeHGNN. SeHGNN (Yang et al., 2023) is one of the earliest and most representative precomputation-based HGNNs, extending SGC (Wu et al., 2019) to heterogeneous graphs. It directly parameterizes spectral filters of the form $S_{r_1}S_{r_2}\cdots S_{r_\ell}$, with $S_r=\tilde{A}_r$ the row-normalized adjacency of relation r, and learns a projection $W_{r_1,\ldots,r_\ell}\in\mathbb{R}^{d\times d'}$ for each such sequence.

Thus, SeHGNN follows Equation 11 exactly, while further applying a Transformer-based semantic fusion followed by an MLP prediction head. The semantic fusion can equivalently be viewed as part of a general decoder $Y = \text{MLP}(\boldsymbol{Z})$, since the metapath sequences (i.e., polynomial terms) form a fixed-length, orderless set.

PSHGCN. PSHGCN (He et al., 2024) is another HGNN that directly learning the coefficients $\theta_0, \theta_{r_1, \dots, r_\ell}$ of SHGC (as in Equation 1) without first-order approximation, while additionally constraining the polynomial coefficients to be positive semi-definite, ensuring the optimization objective convex.

PSHGCN also adopts stronger parameter sharing:

$$W_0 = \theta_0 W, \qquad W_{r_1, \dots, r_\ell} = \theta_{r_1, \dots, r_\ell} W,$$

where $W \in \mathbb{R}^{d \times d'}$ is a node-type-specific projection. Since heterogeneous graphs often have features in different spaces and dimensions, W decomposes as $XW = \sum_{\tau=1}^{|\mathcal{T}|} X_{\tau} W_{\tau}$, where $X = \bigoplus_{\tau=1}^{|\mathcal{T}|} X_{\tau}$ is the direct sum of features across node types and W_{τ} is the corresponding block of W applied to type τ .

C.3 SUMMARY

The above analyses show that representative HGNNs can all be interpreted within the SHGC formulation, differing mainly in (i) the order of approximation and (ii) the degree of parameter sharing. This perspective highlights two key insights: (1) representative models such as RGCN, GTN, and SimpleHGN correspond to first-order SHGC approximations with different parameter-sharing schemes, while SeHGNN and PSHGCN directly instantiate higher-order SHGC; (2) these results provide theoretical justification that atomic HINs are broadly compatible with diverse HGNN designs.

Table 5 provides a concise taxonomy of representative HGNNs under the SHGC lens.

D SRGCN: SIMPLIFIED RELATIONAL GCN

As discussed in Section 3.5, we propose a simplified variant of RGCN under strongest parameter sharing, where relation-specific weights collapse to scalars:

$$W_r^{(\ell)} = \theta_r^{(\ell)} I$$
.

Propagation rule. The update rule of sRGCN is

$$\boldsymbol{H}^{(\ell+1)} = \tilde{\theta}_0^{(\ell)} \boldsymbol{H}^{(\ell)} + \sum_{r=1}^{|\mathcal{R}|} \tilde{\theta}_r^{(\ell)} \tilde{\boldsymbol{A}}_r \boldsymbol{H}^{(\ell)} + \boldsymbol{H}^{(\ell)}, \tag{13}$$

with initial features $\boldsymbol{H}^{(0)} = \boldsymbol{X}\boldsymbol{W}$. Here $\tilde{\theta}_0^{(\ell)}, \{\tilde{\theta}_r^{(\ell)}\}$ are relation coefficients normalized through a softmax, following Yun et al. (2019), for numerical stability.

Multi-head extension. Following Yun et al. (2019); Lv et al. (2021), we employ a multi-head scheme, where multiple sets of coefficients $\{\theta_r^{(\ell,b)}\}_{r,\ell,b}$ are learned in parallel, thereby increasing model capacity.

Training enhancements. We incorporate model-agnostic techniques commonly adopted in advanced HGNNs, including: (i) L2-normalization of output logits (Lv et al., 2021), (ii) edge dropout (or input dropout in the precomputation setting), and (iii) residual connections across layers.

Decoders. Following existing practices, for node classification tasks, we apply an MLP decoder on the learned embeddings. For link prediction tasks, we use either a dot-product decoder or Dist-Mult (Yang et al., 2014).

Precomputation setting. Following the idea of precomputation-based HGNNs (Yang et al., 2023; He et al., 2024), sRGCN can be further simplified by removing intermediate non-linear activations and expanding the layerwise updates into polynomial terms:

$$\boldsymbol{Z} = \left(\prod_{\ell=1}^{L} ((\tilde{\theta}_0 + 1)\boldsymbol{I} + \tilde{\theta}_r^{(\ell)} \tilde{\boldsymbol{A}}_r)\right) \boldsymbol{X} \boldsymbol{W},$$

where $(\tilde{\theta}_0 + 1)$ corresponds to the residual connection.

This enables efficient training on large-scale HINs such as OGBN-MAG, and also permits the use of labels as attributes, as proposed by Yang et al. (2023).

E LIMITATIONS & FUTURE DIRECTIONS

While our study highlights the benefits of *atomic HINs*, several limitations remain. Our treatment of attributes primarily relies on metadata, whereas attributes could in principle be divided into finergrained entities. For example, numerical columns are currently grouped into a single node type, even though each column may carry distinct semantic meaning. Such granularity introduces additional complexity, which remains challenging even for mainstream HGNNs.

Looking ahead, we outline several directions where atomic HINs may enable further advances:

Soft relaxation of selection. Schema refinement in Section 3.2 is framed as binary selections $\alpha_r, \beta_\tau \in \{0,1\}$. These can be relaxed beyond binary: node-type selectors β_τ could interpolate between no identity embedding ($\beta_\tau = 0$) and full-rank embeddings ($\beta_\tau = 1$), with low-rank embeddings corresponding to intermediate values. Similarly, edge-type selectors α_r could be relaxed into continuous meta-weights over relations, particularly suitable for models that parameterize heterogeneity through relation scalars.

Inductive bias of human-defined schemas. Heuristically constructed schemas in common benchmarks often follow practical principles: sparser node types (with larger cardinality) are modeled as structural entities without identity embeddings, while denser node types (with smaller cardinality) are left as attributes. These design choices implicitly define the inductive bias of current HGNNs and may hinder fair evaluation of new models. At the same time, counterexamples highlight the limitations of these heuristics. For instance, the *venue* relation in DBLP and the *subject* relation in ACM provide consistent gains despite their low cardinality, while high-cardinality nodes such as *keyword* or *director* in IMDB also contribute positively when modeled with ID embeddings. These cases underscore the importance of minimizing heuristic assumptions in schema design. More broadly, less conventional structures—such as bottleneck nodes or numerical attributes atomized into dense relation nodes—call for tailored architectural mechanisms in future HGNNs, rather than relying solely on adaptations of existing designs.

LARGE LANGUAGE MODELS

Large Language Models (LLMs) were used for linguistic polishing and editing of the manuscript. All technical content, theoretical results, and experiments were designed and implemented independently by the authors.